arXiv:1403.2439v1 [cs.DM] 10 Mar 2014

String Reconstruction from Substring Compositions

Jayadev Acharya* Hirakendu Das† Olgica Milenkovic‡ Alon Orlitsky§
Shengjun Pan¶

March 12, 2014

Abstract

Motivated by mass-spectrometry protein sequencing, we consider a simply-stated problem of reconstructing a string from the multiset of its substring compositions. We show that all strings of length 7, one less than a prime, or one less than twice a prime, can be reconstructed uniquely up to reversal. For all other lengths we show that reconstruction is not always possible and provide sometimes-tight bounds on the largest number of strings with given substring compositions. The lower bounds are derived by combinatorial arguments and the upper bounds by algebraic considerations that precisely characterize the set of strings with the same substring compositions in terms of the factorization of bivariate polynomials. The problem can be viewed as a combinatorial simplification of the turnpike problem, and its solution may shed light on this long-standing problem as well. Using well known results on transience of multi-dimensional random walks, we also provide a reconstruction algorithm that reconstructs random strings over alphabets of size ≥ 4 in optimal near-quadratic time.

*UCSD. email: jacharya@ucsd.edu
†Yahoo Labs!. email: hdas@yahoo-inc.com
‡UIUC. email: milenkov@uiuc.edu.
§UCSD. email: alon@ucsd.edu.
¶Google. email: s1pan@eng.ucsd.edu.

1 Motivation

A protein is a long sequence of amino acids whose composition and order determine the protein's properties and functionality. A common tool for finding the amino-acid sequence is *mass spectrometry* [1, 2]. It takes a large number of identical proteins, ionizes and randomly breaks them into substrings, and analyzes the resulting mixture to determine the substring weights. The substring weights are then used to infer the amino-acid sequence.

We make two simplifying assumptions that reduce protein mass-spectrometry reconstruction to a simply-stated combinatorial problem that we then analyze.

Assumption 1: The composition of every protein substring can be deduced from its weight. For example, let A, B, and C be three amino acids with respective weights 13, 7, and 4. A string weight of 11 clearly consists of one B and one C. Similarly, weight 18 implies two B's and one C. However, weight 20 could arise from one A and one B or from 5 C's, hence we cannot deduce the composition from the weights. The assumption states that such confusions never arise.

Assumption 2: Each protein-sequence bond gets cut independently with the same probability. For example, if the sequence is ABC and the cut probability is p, then the partition $A|B|C$ is obtained with probability p^2, the partitions $A|BC$, and $AB|C$ are obtained with probability $p(1-p)$, and the partition ABC with probability $(1-p)^2$.

While both assumptions are clearly idealized, one can imagine scenarios where they roughly hold. For example, the first assumption holds if all amino-acid weights are sufficiently large and different. More detailed assumptions, such as some amino-acid weight similarities, or unequal cut probabilities, can be considered as well, see *e.g.,* Section 15, but the two current assumptions provide a clean simple-to-analyze formulation. Additionally, independently of their precise validity, these assumptions convert protein sequencing to a very simple string-reconstruction problem that is interesting on its own merit. It is also a combinatorial simplification of the long-open *turnpike problem* [3, 4], and may provide insight into the structure of its solutions.

The *composition* of a string is the multiset of its elements, namely the number of times each element appears, regardless of the order. Compositions of strings also known as *Parikh vectors*. To derive the reconstruction problem, observe that the first assumption implies that the composition, and hence also the length, of each substring can be determined. The second assumption implies that, ignoring small effects at the sequence ends, all substrings of a given length would appear roughly the same number of times. For example, for cut probability p, strings of length k would appear a number proportional to $p^2(1-p)^{k-1}$. Since the substring lengths can be determined, their number of appearances can be normalized so that the composition of each substring appears exactly once. The problem then is to reconstruct a length-n protein sequence from the multiset of all its $\binom{n+1}{2}$ substring compositions.

For example, the composition of the string $BABCAA$ is the multiset $\{A, A, A, B, B, C\}$, denoted A^3B^2C to indicate that the string consists of three A's, two B's, and one C. Sequencing say the string $ACAB$ would result in $\binom{4+1}{2} = 10$ substring compositions: A, A, B, C, AB, AC, AC, A^2C, ABC, and A^2BC. Note that for each of substring, only the composition, and not the order, is given, and that the compositions are provided along with their multiplicity, but not the location they appear in the string. To reconstruct the original string $ACAB$ from its substring compositions, note that the compositions imply it consists of two A's, a B, and a C, and the two A's don't appear together, hence the string is $ABAC$, $ACAB$, $ABCA$, or their reversals. The appearance of the substring composition A^2C implies $ACAB$ or its reversal.

Clearly a string and its reversal have the same composition multiset and hence cannot be distinguished. We therefore attempt to recover the string up to reversal.

As a final simplification, note that protein sequences are over an alphabet consisting of the 20 amino acids. Yet reconstruction of binary strings extends to reconstructing strings over any finite alphabet. For example, to reconstruct the string $ACAB$ above, first replace appearances

of A's by 1 and of B's and C's by 0. This yields the compositions 0, 0, 1, 1, 01, 01, 01, 0^21, 01^2, 0^21^2, which imply 1010, namely A's appear in the first and third locations. Then replace A's and B's by 1 and C's by 0's to see that B appears in the last location, and finally deduce C's location. A short argument shows that this recursive reconstruction using binary strings always works.

We therefore consider the reconstruction of binary strings from their composition multisets.

2 Definition and results

The *composition multiset* of a string $s = s_1 s_2 \dots s_n \in \{0,1\}^n$ is the multiset

$$\mathcal{S}_s \stackrel{\text{def}}{=} \{\{s_i, s_{i+1}, \dots, s_j\} : 1 \le i \le j \le n\}$$

of compositions of all $\binom{n+1}{2}$ contiguous substrings of s. For example,

$$\mathcal{S}_{001} = \{0, 0, 1, 0^2, 01, 0^21\} \quad \text{and} \quad \mathcal{S}_{010} = \{0, 0, 1, 01, 01, 0^21\}.$$

Note that the number of substrings with any given composition is reflected in the multiset, but their locations are not.

Two strings s and t are *equicomposable*, denoted $s \sim t$, if they have the same composition multisets. Equicomposability is clearly an equivalence relation. The *reversal* of a string $s = s_1 s_2 \dots s_n$ is the string $s^* \stackrel{\text{def}}{=} s_n s_{n-1} \dots s_1$. A string and its reversal are clearly equicomposable, but this is not a problem as in our original motivation they represent the same protein, hence are effectively the same. We say that a string is *reconstructable* if it is equicomposable only with itself and its reversal, hence can be determined up to reversal from its composition multiset. A string s is *non-reconstructable* or *confusable* if it is equicomposable with another string $t \neq s^*$, and if so, we also call s and t *confusable*.

As we did with $ACAB$, it is easy to verify that all binary strings of length at most 7 are reconstructable. A natural question is therefore whether all binary strings are reconstructable. In Section 4 we show that for length 8,

$$01001101 \sim 01101001.$$

Since the two strings are not reversals of each other, they are confusable and neither is reconstructable. An extension of the example yields non-reconstructable binary strings of length n whenever $n+1$ is a product of two integers, each at least 3. For example, of length 8 as $8 + 1 = 3 \cdot 3$, length 11 as $12 = 4 \cdot 3$, etc. That leaves the question as to whether all binary strings of the remaining lengths are reconstructable. Observe that these remaining lengths n are precisely those for which $n+1$ is either 8, a prime, or twice a prime. One of the results we prove is that indeed all strings of these lengths are reconstructable.

To do so, in Sections 5 and 6 we represent a binary string s by a *generating* bivariate polynomial $P_s \in \mathbb{Z}[x, y]$. We also define the *reciprocal* P^* of polynomial P. We show that two strings s and t are equicomposable if and only if

$$P_s P_s^* = P_t P_t^*.$$

We also consider the sets

$$E_s \stackrel{\text{def}}{=} \{t : t \sim s\}$$

and

$$\mathcal{P}_{E_s} \stackrel{\text{def}}{=} \{P_t : t \sim s\}$$

of strings and generating polynomials equicomposable with a string s and show that it has a very simple polynomial characterization. Let

$$P_s = P_1 P_2 \cdots P_k$$

be the prime factorization of P_s over $\mathbb{Z}[x, y]$. It is perhaps intriguing that the equicomposable set of s can be expressed exactly and simply as:

$$\mathcal{P}_{E_s} = \{\tilde{P}_1\tilde{P}_2\cdots\tilde{P}_k : \textit{each } \tilde{P}_i \textit{ is } P_i \textit{ or } P_i^*\}.$$

In Section 9, we use these results to show that indeed if $n+1$ is 8, a prime, or twice a prime, all length-n binary strings are reconstructable. Along with Section 4's construction, this establishes the exact lengths for which all strings are reconstructable.

In Sections 10 and 11 we consider lengths where reconstruction is not always unique. Let

$$E_n \stackrel{\text{def}}{=} \max\{|E_s| : s \in \{0,1\}^n\}$$

be the the largest number of mutually-equicomposable n-bit strings. Since every string is equicomposable with its reversal, $E_n \geq 2$ for every $n \geq 2$, and from the above, $E_n \geq 4$ whenever $n+1$ is a product of two integers ≥ 3.

Generalizing the interleaving construction, we lower-bound E_n, and using the polynomial representation and results on cyclotomic-polynomials, we upper bound it. Considering the distinct-prime factorization $n+1 = 2^{e_0}p_1^{e_1}p_2^{e_2}\ldots p_k^{e_k}$, we show that

$$2^{\lfloor\frac{e_0}{2}\rfloor+e_1+e_2+\cdots+e_k} \leq E_n \leq \min\{2^{(e_0+1)(e_1+1)\cdots(e_k+1)-1}, (n+1)^{1.23}\},$$

where $(e_0+1)(e_1+1)\cdots(e_k+1) = d(n+1)$ is the number of divisors of $n+1$. We also show that when $n+1$ is a prime power or twice a prime power, the lower bound is tight. For all $k \geq 1$,

$$E_{2^k-1} = 2^{\lfloor\frac{k}{2}\rfloor},$$

and for prime $p \geq 3$,

$$E_{p^k-1} = E_{2p^k-1} = 2^k.$$

It follows that for $n = 3^k - 1$,

$$E_n = (n+1)^{\log_3 2}.$$

We then provide a backtracking algorithm that reconstructs a string from it substring compositions. Using some well known results on random walks in high dimensions, we show that the algorithm with high probability reconstructs a string in time $O(n^2 \log n)$ for alphabet size ≥ 4. This is superior, albeit random, compared to the performance of polynomial factorization algorithms. The algorithm can reconstruct strings over alphabet size 2 and 3, although with weaker guarantees.

3 Relation to other work

String reconstruction from multiset decompositions is related to two types of problems. By its formulation it is similar to other string-reconstruction problems, while mathematically it is closer to the well-known turnpike problem.

Several string-reconstruction variations have been previously considered. Reconstruction of a string from a few substrings was considered in [5]. Reconstruction of a string from its subsequences (not necessarily contiguous) was considered in [6–9]. In these problems however, the substrings or subsequences themselves, which include the order of their symbols, are given. By contrast, in our problem, for each substring we are given just the composition, neither the order of the symbols within it nor the substring's location in the original string.

Note that four simple problems of reconstructing a string from its substrings can be formulated. They differ by whether or not we are given (a) the order of the substrings in the string, and (b) the order of the bits in each substring. All four problems are of interest when only some of the $\binom{n+1}{2}$ substrings are given. But when all substrings are given, knowing the order

of either the substrings or their symbols clearly determines the original string. It is only when, as in this mass-spectrometry application, neither the substring order nor the symbol order are provided, that the reconstruction question arises. Some partial results on this problem were presented in [10].

In terms of the solutions and proof techniques, our problem is more closely related to the well-known *turnpike problem* where the locations of n highway exits need to be recovered from the multiset of their $\binom{n}{2}$ inter-exit distances. For example inter-exit distances 1, 2, 3, 3, 5, 6 correspond to exit locations 0, 1, 3, 6.

The turnpike problem originally arose in X-Ray crystallography [11, 12] and has many applications, including to DNA analysis [13]. It is also of theoretical interest as it has an algorithm whose run time is polynomial in the largest inter-exit distance, but it is not known whether it has an algorithm whose run time is polynomial in n, independently of the distances. Several variation of the problem have been recently considered, e.g. [14, 15] and references therein. A related problem of characterizing strings that have the same *set* of compositions, instead of multisets, is considered in [16].

String reconstruction from substring multisets is a combinatorial simplification of the turnpike problem [15, 17]. To see that, convert every string-reconstruction instance to a turnpike problem whose solution implies the original string. Given an n-bit string reconstruction problem, replace each substring composition by an inter-exit distance obtained by replacing each 0 by 1 and each 1 by $n+1$ and summing the values. Then find the exit locations and replace back the bit values.

For example, to recover the string 1011 from the compositions 0, 1, 1, 1, 01, 01, 11, 01^2, 01^2, 01^3, replace every 0 by 1 and every 1 by $4+1=5$ to obtain inter-exit distances 1, 5, 5, 5, 6, 6, 10, 11, 11, 16 which in turn correspond to locations 0, 5, 6, 11, 16. Then convert 1 back to 0, and 5 back to 1 to obtain the original string 1011. Note that the mapping of 1 to $n+1$ prevents spurious solutions.

This reduction shows that string reconstruction can be solved in polynomial time [18], but this is not our main focus. Instead, we show that several questions unsolved for the turnpike problem can be answered for string reconstruction. Hence, in addition to the problem's intrinsic value, and potential application to protein reconstruction, its solution may provide useful insights for the general turnpike problem.

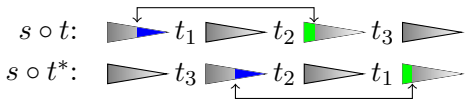
For example, some of the more important unsolved questions about the turnpike problem concern the number of solutions a given instance may have [4]. A turnpike problem with $n \geq 6$ exits may have multiple solutions, but their largest possible number is not known for any such n. By contrast, for string reconstruction we show that when $n+1$ is 8, a prime, or twice a prime reconstruction is always unique. We also determine the exact number of maximal reconstructions whenever $n+1$ is a prime power or twice a prime power, and conjecture that the same formula holds for all string lengths.

We also note that some of the polynomial techniques we use are related to those applied to the turnpike problems [4, 19]. Yet others, such as the bivariate polynomial formulation and relation to cylcotomic polynomials, seem new.

4 Simple confusions

A simple computer search shows that all binary strings of length at most 7 are reconstructable. For length 8, the strings 01001101 and 01101001 are confusable and are not reversals of each other. To see that, note that they can be parsed as **01** 0 **01** 1 **01** and **01** 1 **01** 0 **01**. Both have a common substring **01**, which is interleaved with 0 1 for the first string and with its reversal 1 0 in the second string. The *interleaving* of string s with the bits of $t = t_1 \dots t_m$ is the string $s \circ t \stackrel{\text{def}}{=} s\, t_1\, s\, t_2 \cdots t_m\, s$. The first string above is therefore $01 \circ 01$ while the second is $01 \circ 10$.

Recall that t^* represents the reversal of t, and that $\sim$ indicates that two strings are equicomposable. We show that $s \circ t \sim s \circ t^*$ for any s and t. We demonstrate that fact when t has length 3, and the general result follows similarly. The interleaving $s \circ t$ is represented by the figure below, where each triangle represents the string s.



Clearly any substring of $s \circ t$ can be uniquely mapped to a substring of $s \circ t^*$ with the same composition. For example, as illustrated in the figure, the substring in $s \circ t$ consisting of a tail of s (blue), t_1, s, t_2, and a head of s (green) has the same composition as the substring in $s \circ t^*$ consisting of the same tail of s (blue), t_2, s, t_1, and the same head of s (green). Thus $s \circ t$ and $s \circ t^*$ have the same multiset of compositions, and hence are equicomposable.

If the length of s and t are m and m' respectively, then the length of $s \circ t$ is $n = (m+1)(m'+1) - 1$. If $m \geq 2$, $m' \geq 2$, we can always choose non-palindromic t, namely $t \neq t^*$, ensuring that $s \circ t \neq s \circ t^*$, and non-palindromic s, namely $s \neq s^*$, ensuring that $s \circ t \neq (s \circ t^*)^*$. It follows that whenever $n+1$ is a product of two integers, each ≥ 3, there are confusable n-bit strings.

5 Polynomial representation

The previous section described some confusable strings. To further characterize confusability, we represent strings as polynomials. A similar representation has been used for the turnpike problem [4, 19], however the polynomials used there are univariate, whereas bivariate polynomials are better suited for string reconstruction.

We use this representation to show that the strings equicomposable with a string s can be determined by the prime factorization of the polynomial representing s, and that s can be reconstructed from its composition multiset by polynomial factorization.

All polynomials and factorizations in the paper are over $\mathbb{Z}[x, y]$. A polynomial whose nonzero coefficients are all 1 is *0-1*. A term $x^a y^b$ is a *monomial* with *x-degree* a, *y-degree* b, and *total degree* $a+b$. The *x-*, *y-*, and *total-degrees* of a polynomial are the corresponding highest degrees of any of its monomials. For example, for $x^2y + xy^3$ they are 2, 3, and 4 respectively.

When representing strings by polynomials, 0 is denoted by x and 1 by y. Also, a_i denotes the number of 0's in the first i bits of a string and a denotes their total number, while b_i and b represent the same for the 1's. For example, 01011 has $a_3 = a = 2$, $b_3 = 1$, and $b = 3$.

The *generating polynomial* of a binary string $s = s_1 s_2 \cdots s_n$ is

$$P_s(x, y) \stackrel{\text{def}}{=} \sum_{i=0}^{n} x^{a_i} y^{b_i} = \sum_{i=0}^{n} x^{a_i} y^{i-a_i}.$$

For example,

$$P_{0100}(x, y) = 1 + x + xy + x^2y + x^3y.$$

Generating polynomials of n-bit strings are characterized by the following sufficient properties.

G1 They are 0-1.

G2 They have $n+1$ terms, exactly one of each total degree $0, 1, \ldots, n$.

G3 For all $1 \leq i \leq n$, the ratio between the terms of total degrees i and $i-1$ is either x or y.

For example, $1 + x + xy + x^2y + x^3y$ is generating, but $1 - x$, $1 + x^2$, and $1 + x + y^2$ are not.

Similar to strings, a composition 0^a1^b is represented by the monomial x^ay^b, and $\mathcal{S}_s(x, y)$ is the summation of all monomials corresponding to the substring compositions of s. For example,

$$\mathcal{S}_{0100}(x, y) = \{0, 0, 0, 1, 0^2, 01, 01, 0^21, 0^21, 0^31\} = 3x + y + x^2 + 2xy + 2x^2y + x^3y.$$

Note that for all $i \leq j$,

$$x^{a_j}y^{b_j} \cdot (x^{a_i}y^{b_i})^{-1} = x^{a_j-a_i}y^{b_j-b_i}$$

represents the composition of the substring $s_{i+1} \dots s_j$, hence the generating polynomial provides a simple expression for the composition multiset,

$$P_s(x, y)P_s\left(\frac{1}{x}, \frac{1}{y}\right) = n + 1 + \mathcal{S}_s(x, y) + \mathcal{S}_s\left(\frac{1}{x}, \frac{1}{y}\right).$$

For instance, in our running example,

$$\begin{aligned} P_{0100}(x, y)P_{0100}\left(\frac{1}{x}, \frac{1}{y}\right) &= 5 + 3x + y + x^2 + 2xy + 2x^2y + x^3y + \frac{3}{x} + \frac{1}{y} + \frac{1}{x^2} + \frac{2}{xy} + \frac{2}{x^2y} + \frac{1}{x^3y} \\ &= 5 + \mathcal{S}_{0100}(x, y) + \mathcal{S}_{0100}\left(\frac{1}{x}, \frac{1}{y}\right). \end{aligned}$$

To reconstruct a string from a composition multiset $\mathcal{S}$, we therefore find all strings s whose generating polynomial P_s satisfies the above equation. We also have the following basic lemma.

Lemma 1. *Strings s and t are equicomposable iff*

$$P_s(x, y)P_s\left(\frac{1}{x}, \frac{1}{y}\right) = P_t(x, y)P_t\left(\frac{1}{x}, \frac{1}{y}\right). \tag{1}$$

We next show that for a string s, an equation like (1) holds only for generating polynomials. The approach is similar to [4, 19] with additional work needed to prove that the polynomial obtained satisfied all the properties of generating polynomials.

Lemma 2. *If P is a generating polynomial and $Q \in \mathbb{Z}[x, y]$ with $Q(0, 0) > 0$ and*

$$P(x, y)P\left(\frac{1}{x}, \frac{1}{y}\right) = Q(x, y)Q\left(\frac{1}{x}, \frac{1}{y}\right), \tag{2}$$

then $Q(x, y)$is a generating polynomial.

Proof. We first show that Q satisfies **G1**.
G1 Evaluating (2) at $x = y = 1$, we obtain $P(1, 1)^2 = Q(1, 1)^2$, namely $P(1, 1) = \pm Q(1, 1)$. Suppose '+'. Let $P(x, y) = \sum_{i,j} p_{i,j}x^iy^j$ and $Q(x, y) = \sum_{i,j} q_{i,j}x^iy^j$, where $p_{i,j}, q_{i,j} \neq 0$, then,

$$\sum_{i,j} p_{i,j} = \sum_{i,j} q_{i,j}.$$

Comparing the constant term in (2),

$$\sum_{i,j} p_{i,j}^2 = \sum_{i,j} q_{i,j}^2.$$

Subtracting the two equations,

$$\sum_{i,j} p_{i,j}(1 - p_{i,j}) = \sum_{i,j} q_{i,j}(1 - q_{i,j}).$$

Since all $p_{i,j}$ are 1, the left and hence the right sides are 0. For all integers i, $i(1-i) \leq 0$ with equality iff i is either 0 or 1. Since all $q_{i,j}$ are nonzero integers, they must be 1. The case of '$-$' is similar.

G2 Each term in $P(x,y)P(1/x,1/y)$ is of the form $x^{a_j-a_i}y^{b_j-b_i}$. Thus the exponents of x and y cannot have different signs. Q cannot have two terms of the same total degree as their corresponding product in $Q(x,y)Q(1/x,1/y)$ would yield a term with x- and y-degrees of opposing signs, which by G1 would not get cancelled by any other product term. Let $0 = d_0 < d_1 < \ldots < d_n$ be the degrees of the terms in Q. The largest degree term in (2) on the LHS is $x^{a_n}y^{n-a_n}$ whose degree is n. The largest degree on the RHS is $d_n - d_0 = d_n$. So $d_i = d_0 + i = i$ and property G2 is proved.

G3 Consider the linear terms (constant times x or y) in (2). For any 0-1 polynomial $F(x,y)$, the sum of the coefficients of the linear terms in $F(x,y)F(1/x,1/y)$ is the number of pairs of terms in F whose ratio is x or y. P has n such pairs, hence so does Q. By G2, Q has $n+1$ terms one of each total degree $0, 1, \ldots, n$, the ratio between any two consecutive terms must be x or y. □

The two lemmas characterize the generating polynomials of all the strings in E_s.

Theorem 3.

$$\mathcal{P}_{E_s} = \{P(x,y) \in \mathbb{Z}[x,y] : P(0,0) > 0, P(x,y)P\left(\frac{1}{x},\frac{1}{y}\right) = P_s(x,y)P_sP\left(\frac{1}{x},\frac{1}{y}\right)\}$$

6 Reciprocal polynomials

To apply existing results on polynomial factorization, we relate $P(1/x,1/y)$ to standard polynomials. Let $\deg_x P$ and $\deg_y P$ be the highest x- and y-degrees of a bivariate polynomial P. The *reciprocal* of P is the polynomial

$$P^*(x,y) \stackrel{\text{def}}{=} x^{\deg_x P}y^{\deg_y P}P\left(\frac{1}{x},\frac{1}{y}\right).$$

For example,

$$(y + 3xy - 2y^2)^* = xy^2\left(\frac{1}{y} + \frac{3}{xy} - \frac{2}{y^2}\right) = xy + 3y - 2x.$$

We will use the following reciprocal-polynomial properties.

R1 $\deg_x P^* \leq \deg_x P$ with equality iff $x \nmid P$, and similarly for y.

R2 The reciprocal of the product is the product of the reciprocals: $(P_1P_2)^* = P_1^*P_2^*$.

R3 The reciprocal of the generating polynomial of s generates the reversal of s: $P_s^* = P_{s^*}$.

For example,

$$P_{0100}^*(x,y) = (1 + x + xy + x^2y + x^3y)^* = 1 + x + x^2 + x^2y + x^3y = P_{0010}(x,y).$$

These properties imply a polynomial formulation of Lemmas 1 and 2 and Theorem 3. The proofs are almost identical to the proofs of their counterparts shown earlier and are hence omitted.

Lemma 4. *Strings s and t are equicomposable iff*

$$P_sP_s^* = P_tP_t^*.$$

The second lemma has actually a slightly cleaner formulation than its $1/x$ counterpart.

Lemma 5. *If P is a generating polynomial and for $Q \in \mathbb{Z}[x, y]$*

$$PP^* = QQ^*,$$

then Q or its negation is also a generating polynomial.

Theorem 6.

$$\mathcal{P}_{E_s} = \{P(x, y) \in \mathbb{Z}[x, y] : P(x, y) > 0, P(x, y)P^*(x, y) = P_s(x, y)P_s^*(x, y)\}$$

So far, we studied products of generating polynomials and their reciprocals. The next lemma, addresses the factors of generating polynomials.

Lemma 7. *$s \sim t$ if and only if for some $A, B \in \mathbb{Z}[x, y]$,*

$$P_s = AB \qquad \text{and} \qquad P_t = AB^*.$$

Proof. Let $A = \gcd(P_s, P_t)$, then $P_s = AB$ and $P_t = AC$ for relatively prime B and C. We show that $C = B^*$.

By Lemma 4, $P_sP_s^* = P_tP_t^*$. By R2, $P_s^* = A^*B^*$ and $P_t^* = A^*C^*$, hence, $BB^* = CC^*$. Since B and C are relatively prime, C divides B^*. But since P_s and P_t have the same total degree, B, B^*, C, and C^*, all have the same total degree. And since P_s and P_t are generating polynomials, the four polynomials have constant term 1. Hence $C = B^*$.

Conversely if $P_s = AB$ and $P_t = AB^*$ then it is easy to see that $P_sP_s^* = (AB)(AB)^* = AA^*BB^* = AB^*(AB^*)^* = P_tP_t^*$. □

For example, we saw that the 8-bit strings $s = 01001101$ and $t = 01101001$ are confusable. Indeed $P_s(x, y) = (1 + x + xy)(1 + x^2y + x^3y^3)$ and $P_t(x, y) = (1 + x + xy)(1 + xy^2 + x^3y^3) = (1 + x + xy)(1 + x^2y + x^3y^3)^*$. Here, A is a generating polynomial, but this is not always the case. The 23-bit strings 01000101010000100011001 and 01010100010000110010001 are also confusable, but here, $A = 1+y+xy+xy^2+x^4y^8+x^4y^9+x^4y^{10}+x^5y^{10}$ and $B = 1+xy^3+x^3y^5$. In the above examples, A and B are 0-1 polynomials, but that too is not always the case. The string 01001001001 can be decomposed as $(1 + x + x^3y - x^4y^2 + x^3y^3 + x^5y^3 + x^5y^4)(1 + xy)$ where the first factor has a negative coefficient.We will return to the structure of these factors in Section 10.

The previous results imply a very simple expression for E_s, the set of strings equicomposable with s. Consider the factorization of P_s into irreducible factors (prime factorization)

$$P_s = P_1P_2 \cdots P_k$$

The next theorem says that every product of the P_i's or their reciprocals is a generating polynomial and these generating polynomials exactly correspond to all strings equicomposable with s.

Theorem 8. *For every s,*

$$\mathcal{P}_{E_s} = \{\tilde{P}_1\tilde{P}_2 \cdots \tilde{P}_k : \textit{each } \tilde{P}_i \textit{ is } P_i \textit{ or } P_i^*\}.$$

Proof. If $P = \tilde{P}_1\tilde{P}_2 \cdots \tilde{P}_k$, where each $\tilde{P}_i$ is either P_i or P_i^*, let A be the product of the unmodified P_i's and B be the product of the P_i's that are reciprocated. Then $P_s = AB$ and $P = AB^*$. Hence $P_sP_s^* = PP^*$, and by Lemma 5, $P = P_t$ for some string t. But then $P_sP_s^* = P_tP_t^*$, and by Lemma 4, $t \in E_s$.

Conversely, if $t \in E_s$, then by Lemma 7, $P_s = AB$ while $P_t = AB^*$. The prime factorizations $A = P_1P_2 \cdots P_j$ and $B = P_{j+1} \cdots P_k$ yield the prime factorizations $P_s = P_1P_2 \cdots P_k$ and $P_t = P_1P_2 \cdots P_jP_{j+1}^* \cdots P_k^*$, hence P_t can be written as $\tilde{P}_1\tilde{P}_2 \cdots \tilde{P}_k$. □

The theorem provides a simple formula for the number of strings equicomposable with any given string s. A polynomial is *palindromic* if it equals its reciprocal. For example, $1 + x + xy + xy^2 + x^2y^2$, but not $1 + x + xy$ ($(1 + x + xy)^* = 1 + y + xy$). Let ν_s be the number of non-palindromic terms in the prime factorization of $P_s(x, y)$.

Corollary 9. *For every string s,*

$$|E_s| \leq 2^{\nu_s}.$$

In particular, if $\nu_s = 1$, then s is reconstructable.

Proof. The first part follows trivially from the theorem. The second part follows as $\nu_s = 0$ implies $E_s = \{s\}$, and $\nu_s = 1$ implies $E_s = \{s, s^*\}$. □

In Section 8 we will return to this result and show that the inequality holds with equality.

7 Cyclotomic polynomials

Theorem 8 characterizes equicomposability in terms of prime factorization of generating polynomials. Factoring bivariate polynomials is hard in general, but the current analysis is simplified by the fact that any factorization of a generating polynomial $P(x, y)$ implies a factorization of the univariate polynomial $P(x, x)$ obtained by evaluating $P(x, y)$ at $y = x$. From properties G1 and G2 of generating polynomials, if $P(x, y)$ generates an n-bit string, then

$$P(x, x) = 1 + x + x^2 + \cdots + x^n = \frac{x^{n+1} - 1}{x - 1}.$$

Factorizations of $x^n - 1$, and hence of $x^{n+1} - 1$, have been studied extensively and for completeness we present a small part of the literature that we use to analyze confusability.

The n complex numbers $e^{\frac{2\pi i}{n}}, e^{2\frac{2\pi i}{n}}, \ldots, e^{n\frac{2\pi i}{n}} = 1$, whose nth powers are 1, are the nth *roots of unity*. They are exactly the roots of $x^n - 1$, and therefore

$$x^n - 1 = \prod_{\omega}(x - \omega),$$

where the product ranges over all nth roots of unity. The prime factorization of $x^n - 1$ over $\mathbb{Z}[x]$ therefore partitions these n terms into groups, each multiplying to an irreducible polynomial over $\mathbb{Z}[x]$.

For a positive integer d, a dth root of unity is *primitive* if none of its smaller positive powers is 1. For example:

d	1	2	3	4	5	6
primitive dth roots of unity	1	-1	$e^{\pm 2\pi i/3}$	$\pm i$	$e^{2\pi ij/5}, j = 1, \ldots, 4$	$e^{\pm 2\pi i/6}$

The dth *cyclotomic polynomial* is

$$\Phi_d(x) \stackrel{\text{def}}{=} \prod_{\omega}(x - \omega)$$

where the product ranges over all primitive dth roots of unity. For example,

$$\begin{aligned}
\Phi_1(x) &= x - 1 \\
\Phi_2(x) &= x + 1 \\
\Phi_3(x) &= (x - e^{2\pi i/3})(x - e^{-2\pi i/3}) = x^2 + x + 1 \\
\Phi_4(x) &= (x - i)(x + i) = x^2 + 1 \\
\Phi_5(x) &= \prod_{j=1}^{4}(x - e^{\frac{2\pi ij}{5}}) = x^4 + x^3 + x^2 + x + 1 \\
\Phi_6(x) &= (x - e^{2\pi i/6})(x - e^{-2\pi i/6}) = x^2 - x + 1.
\end{aligned}$$

While we will not use this fact, note that the degree of $\Phi_d(x)$ is the number of integers between 1 and d that are relatively prime with d.

Every nth root of unity is a primitive dth root of unity for some $1 \le d \le n$, and it is easy to see that $d \mid n$. Hence, for every n,

$$x^n - 1 = \prod_{d|n} \Phi_d(x). \tag{3}$$

For example,

$$\begin{aligned}
x^2 - 1 &= (x-1)(x+1) \\
x^3 - 1 &= (x-1)(x^2+x+1) \\
x^4 - 1 &= (x-1)(x+1)(x^2+1) \\
x^5 - 1 &= (x-1)(x^4+x^3+x^2+1) \\
x^6 - 1 &= (x-1)(x+1)(x^2+x+1)(x^2-x+1).
\end{aligned}$$

It is easy to see that Cyclotomic polynomials have integer coefficients. Comparing coefficients shows that if $a(x), b(x) \in \mathbb{Q}[x]$ and $a(x)b(x) \in \mathbb{Z}[x]$, then $a(x), b(x) \in \mathbb{Z}[x]$. By induction on n, assume that for all $d < n$ with $d \mid n$, $\Phi_d(x) \in \mathbb{Z}[x]$. Then (3) implies that

$$\Phi_n(x) = \frac{x^n - 1}{\prod_{d<n, d|n} \Phi_d(x),}$$

and since different cyclotomic polynomials are relatively prime (no two cyclotomic polynomials share a root), $\Phi_n(x) \in \mathbb{Q}[x]$ and by the fact mentioned above, is $\in \mathbb{Z}[x]$. Gauss showed, *e.g.,* [20], that the cyclotomic polynomials are irreducible. The irreducibility proof of general cyclotomic polynomials is somewhat involved. However, we will mostly need to factor $x^p - 1$ and $x^{2p} - 1$ for a prime p. Their factorization require the irreducibility of $\Phi_p(x)$ and $\Phi_{2p}(x)$, which follow easily from *Eisenstein's Criterion.*

Lemma 10 (Eisenstein's Criterion)**.** *Let $P(x) = a_n x^n + a_{n-1}x^{n-1} + \cdots + a_0$. If some prime $p \mid a_0, a_1, \ldots, a_{n-1}$ but $p \nmid a_n$ and $p^2 \nmid a_0$, then $P(x)$ is irreducible over $\mathbb{Z}$.*

Proof. Suppose towards a contradiction that $P(x) = (b_k x^k + \cdots + b_0)(c_\ell x^\ell + \cdots + c_0)$ where $k, \ell \ge 1$ and $b_k, c_\ell \ne 0$. Since p divides $a_0 = b_0 c_0$ but p^2 does not, p divides exactly one of b_0 and c_0, without loss of generality, b_0. Since p does not divide $a_n = b_k c_\ell$, there is a smallest index $i < k < n$ such that p does not divide b_i. It follows that $p \mid a_i, b_0, b_1, \ldots, b_{i-1}$ and $p \nmid c_0, b_i$, contradicting $a_i = b_i c_0 + b_{i-1}c_1 + \cdots + b_0 c_i$. □

Eisenstein Criterion's best-known application is to the irreducibility of the cyclotomic polynomial $\Phi_p(x)$ for a prime p.

Corollary 11. *For every prime p,*

$$\Phi_p(x) = 1 + x + x^2 + \cdots + x^{p-1}$$

is irreducible over $\mathbb{Z}$.

Proof. Substituting x by $x + 1$,

$$\Phi_p(x+1) = \frac{(x+1)^p - 1}{x} = \binom{p}{1} + \binom{p}{2}x + \cdots + \binom{p}{p}x^{p-1}.$$

Since p is prime, it divides $\binom{p}{1}, \binom{p}{2}, \ldots, \binom{p}{p-1}$, but not $\binom{p}{p}$. Therefore p divides all coefficients except the leading. Furthermore, p^2 does not divide the constant coefficient $\binom{p}{1}$. By Eisenstein's Criterion, $\Phi_p(x+1)$ and therefore $\Phi_p(x)$ are irreducible over $\mathbb{Z}$. □

Note that $\Phi_{2p}(x) = \Phi_p(-x)$ for odd primes p, hence the irreducibility of $\Phi_p(x)$ suffices for our purposes.

Lemma 12. *$\Phi_d(x)$ is palindromic for all $d \geq 2$.*

Proof. For any polynomial $P(x)$, the roots of $P^*(x)$ over $\mathbb{C}$ are exactly the reciprocals of the roots of $P(x)$, and the constant coefficient of $P^*(x)$ is the leading coefficient of $P(x)$. Since the reciprocal polynomial always exists, and the roots and constant coefficient specify a polynomial, these two conditions are also sufficient for a polynomial to be the reciprocal of $P(x)$.

For every d, the roots of $\Phi_d(x)$ are on the unit circle, hence the reciprocal of a root is also its complex conjugate. Furthermore, since $\Phi_d(x)$ has real coefficients, every root appears with its complex conjugate, hence the reciprocals of the roots are roots too. Finally, for $d \geq 2$, $\Phi_d(x)$ has the same leading and constant coefficient, namely 1, hence it is its own reciprocal. □

8 The size of confusable sets

The next lemma shows that in any factorization of a generating polynomial, no two terms are reciprocal of each other. Our proof is different and simpler than the one used in [4] for the general turnpike problem.

Lemma 13. *Generating polynomials do not have two mutually reciprocal factors.*

Proof. Let $P(x, y)$ generate an n-bit string. If $A(x, y)A^*(x, y)$ divides $P(x, y)$, then $A(x, x)A^*(x, x)$ $= A^2(x, x)$ divides $P(x, x) = 1 + x + x^2 + \cdots + x^n$. Clearly, $A^2(x, x)$ has double roots over $\mathbb{C}$. The roots of $1 + x + \cdots + x^n$ are all distinct, hence it has no multiple roots. □

If follows that Corollary 9 holds with equality.

Corollary 14. *For every string s,*

$$|E_s| = 2^{\nu_s},$$

and s is reconstructable iff $\nu_s \leq 1$. □

For example, for $s = 01001101$ the factorization of its generating polynomial into irreducible factors is $P_s(x, y) = (1 + x + xy)(1 + x^2y + x^3y^3)$, hence $\nu_s = 2$. The four confusable strings are $E_{01001101} = \{01001101, 01101001, 10110010, 10010110\}$.

9 Unique reconstruction for prime-related lengths

In Section 4, we saw that whenever $n + 1$ is a product of two integers ≥ 3, there are confusable n-bit strings. The remaining lengths are those where $n + 1$ is a prime, twice a prime, or 8. We use the polynomial representation to show that for all these lengths reconstruction is always unique.

While factorization and irreducibility of $P(x, y)$ is related to that of $P(x, x)$, the two are not in complete correspondence. Irreducibility of $P(x, y)$ does not imply irreducibility of $P(x, x)$. For example, $P(x, y) = 2x^2 - y^2 = (\sqrt{2}x + y)(\sqrt{2}x - y)$ hence is irreducible over $\mathbb{Z}$, yet $P(x, x) = x^2$ is reducible.

Similarly, in the direction we need, irreducibility of $P(x, x)$ does not generally imply irreducibility of $P(x, y)$. For example, $P(x, y) = (x - y + 1)(x + 1)$ is reducible even though $P(x, x) = x + 1$ is irreducible. Yet the next two lemmas show that for generating polynomials, and more generally when $\deg P(x, y) = \deg P(x, x)$, this implication holds.

Lemma 15. *Any factor of a generating polynomial has a single term of highest total degree.*

Proof. Let $\widehat{P}(x,y)$ be the sum of the highest-total-degree terms in a polynomial $P(x,y)$. For example, if $P(x,y) = x^3 + x^2y^2 + y^4$ then $\widehat{P}(x,y) = x^2y^2 + y^4$. If $A(x,y)$ is a factor of $P(x,y)$, then $\widehat{A}(x,y)$ divides $\widehat{P}(x,y)$, and if $P(x,y)$ is also generating, then $\widehat{P}(x,y)$ consists of a single term, hence so does its factor $\widehat{A}(x,y)$. □

Lemma 16. *The number of terms in the prime factorization of a generating polynomial $P(x,y)$ is at most the corresponding number for $P(x,x)$. Hence if $P(x,x)$ is irreducible, so is $P(x,y)$.*

Proof. Let $A_1(x,y)A_2(x,y)\cdots A_k(x,y)$ be the prime factorization of $P(x,y)$. Then $A_1(x,x)$ $A_2(x,x)\cdots A_k(x,x)$ is a factorization of $P(x,x)$ into polynomials that by Lemma 15 are non-constant. □

Theorem 17. *All strings of length 7 are reconstructable.*

Proof. A simple computer search proves the result, but we also provide a theoretical proof in Appendix A. □

Theorem 18. *All strings of length one less than a prime are reconstructable.*

Proof. If s is $(p-1)$-bits long then by the properties of generating polynomials,

$$P_s(x,x) = 1 + x + x^2 + \ldots + x^{p-1}.$$

By Corollary 11, $P_s(x,x)$ is irreducible. By Lemma 16, $P_s(x,y)$ is irreducible. By Corollary 9, s is reconstructable. □

For example, for $n = 2$, both $P_{00}(x,y) = 1 + x + x^2$ and $P_{01}(x,y) = 1 + x + xy$, and their symmetric versions, $P_{10}(x,y)$ and $P_{11}(x,y)$, are irreducible.

To prove the result for twice a prime, we'll need the following simple lemma.

Lemma 19. *If s is confusable, then $P_s = AB$ for some non-palindromic polynomials A and B.*

Proof. Let s be confusable with t. By Lemma 7, $P_s = AB$ and $P_t = AB^*$ for some A and B. If B is palindromic, then $P_s = AB = AB^* = P_t$, namely $s = t$, while if A is palindromic, then $P_s^* = A^*B^* = AB^* = P_t$, namely $s = t^*$. □

Theorem 20. *All strings of length one less than twice a prime are reconstructable.*

Proof. Let s have length one less than twice a prime. We show that in any factorization

$$P_s(x,y) = f(x,y)g(x,y),$$

at least one of $f(x,y)$ and $g(x,y)$ is palindromic. By Lemma 19, s is reconstructable.

Eisenstein's criterion again implies the prime factorization

$$\begin{aligned} P_s(x,x) &= 1 + x + x^2 + \cdots + x^{2p-1} \\ &= (1+x)(1 - x + x^2 - \cdots + x^{p-1})(1 + x + \cdots + x^{p-1}). \end{aligned}$$

Hence there are only three factorizations of $P_s(x,y)$ into two factors.

Case 1: $f(x,x) = 1 + x$ and $g(x,x) = (1 - x + x^2 - \cdots + x^{p-1})(1 + x + \cdots + x^{p-1})$,

Case 2: $f(x,x) = 1 + x + \cdots + x^{p-1}$ and $g(x,x) = (1+x)(1 - x + x^2 - \cdots + x^{p-1}) = 1 + x^p$,

Case 3: $f(x,x) = 1 - x + \cdots + x^{p-1}$ and $g(x,x) = (1+x)(1 + x + \cdots + x^{p-1})$.

We show that in all three cases, at least one of $f(x,y)$ and $g(x,y)$ is palindromic.

Case 1: By Lemma 15, $f(x,y)$ must be either $1 + x$ or $1 + y$, and both are palindromic.

Case 2: Let $s = s_1 s_2 \dots s_n$. Then $P(x,y) \stackrel{\text{def}}{=} P_s(x,y) = \sum_{i=0}^{2p-1} x^{a_i} y^{i-a_i}$. Denote $\frac{\partial P}{\partial x}(x,y)|_{y=x}$ by P'_x to obtain

$$P'_x = \sum_{i=1}^{2p-1} a_i x^{i-1}.$$

Therefore,

$$\frac{\partial P}{\partial x}(x,y) = f(x,y)\frac{\partial g}{\partial x}(x,y) + g(x,y)\frac{\partial f}{\partial x}(x,y)$$

implies

$$P'_x = (1 + x + \cdots + x^{p-1})g'_x + (1+x^p)f'_x,$$

and multiplying both sides by $1-x$,

$$(1-x)P'_x = (1-x^p)g'_x + [x^p(1-x) + (1-x)]f'_x.$$

The left-hand side can be written as

$$\sum_{i=1}^{2p-1} a_i x^{i-1} - \sum_{i=1}^{2p-1} a_i x^i = a_1 + \sum_{i=1}^{2p-2}(a_{i+1} - a_i)x^i - a_{2p-1}x^{2p-1} = \sum_{i=0}^{2p-2} s_{i+1}x^i - a_{2p-1}x^{2p-1}$$

and the right-hand side as

$$\underbrace{\left[(1-x)f'_x + g'_x\right]}_{(I)} + \underbrace{x^p \cdot \left[(1-x)f'_x - g'_x\right]}_{(II)}.$$

Note that $\deg(I) \le \max\{\deg f, \deg g - 1\} = p-1$, while all terms in (II) have degree $\ge p$. Hence,

$$(1-x)f'_x + g'_x = \sum_{i=0}^{p-1} s_{i+1}x^i,$$
$$(1-x)f'_x - g'_x = \sum_{i=p}^{2p-2} s_{i+1}x^{i-p} - a_{2p-1}x^{p-1} = \sum_{i=0}^{p-2} s_{i+1+p}x^i - a_{2p-1}x^{p-1}.$$

Subtracting the two equations,

$$2g'_x = \sum_{i=0}^{p-2}(s_{i+1} - s_{i+1+p})x^i + (s_p + a_{2p-1})x^{p-1}.$$

Since g has integer coefficients and the s_i's are bits, for $i = 1, 2 \dots, p-1$,

$$s_i = s_{i+p},$$

and hence $s = tut = t \circ u$, where t consists of $p-1$ bits and u is a single bit. Letting a and b denote the number of zeros and ones in tu, $P_s(x,y) = (1 + x^a y^b)P_t(x,y)$. By theorem 18 $P_t(x,y)$ has at most one non-palindromic factor. It suffices to show that $1 + x^a y^b$ which itself is palindromic has no non-palindromic factors. If $a = 0$ or $b = 0$, $1 + x^a y^b$ has no non-palindromic factor. If both a and b are non-zero, since $a + b = p$, it has at most two factors one of which by Lemma 15 must be $1 + x$ or $1 + y$ which is not possible since $x = -1$ or $y = -1$ is not a root of $1 + x^a y^b$.

Case 3: The proof follows as in Case 2 with modulo 2 operations,

$$g(x,x) = (1+x)(1 + x + \cdots + x^{p-1}) = 1 + x^p \pmod 2.$$ □

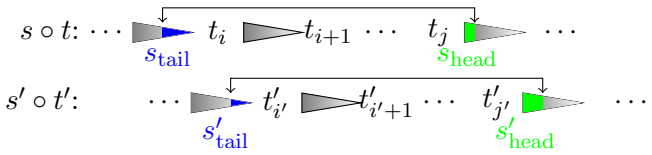
10 Lower bound on E_n

Recall that E_n is the largest number of mutually-equicomposable n-bit strings. In Section 4 we saw that for all strings s, t, the interleaved strings $s \circ t$ and $s \circ t^*$ are equicomposable. Choosing non-palindromic s and t, we derived confusable strings for all lengths n such that $n+1$ is a product of two integers ≥ 3. The interleaved strings are non-palindromic, hence $E_n \geq 4$. We now generalize the interleaved construction in two ways to construct equicomposable sets that for some n are as large as $(n+1)^{\log_3 2}$.

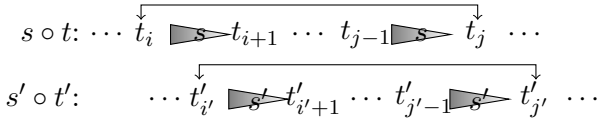
We first show that if $s \sim s'$ and $t \sim t'$, then

$$s \circ t \sim s' \circ t'.$$

Any substring of $s \circ t$ is of the form $s_{\text{tail}}\, t_i\, s\, t_{i+1}\, s \cdots s\, t_j\, s_{\text{head}}$, where the *tail* and *head* substrings $s_{\text{tail}} = s_\ell s_{\ell+1} \cdots s_n$ and $s_{\text{head}} = s_1 s_2 \cdots s_h$ can also be empty or all of s. We show that it can be bijectively mapped to a substring $s'_{\text{tail}}\, t'_{i'}\, s'\, t'_{i'+1}\, s' \cdots s'\, t'_{j'}\, s'_{\text{head}}$ of $s' \circ t'$, where $s'_{\text{tail}} = s'_{\ell'} s'_{\ell'+1} \cdots s'_n$ and $s'_{\text{head}} = s'_1 s'_2 \cdots s'_{h'}$.

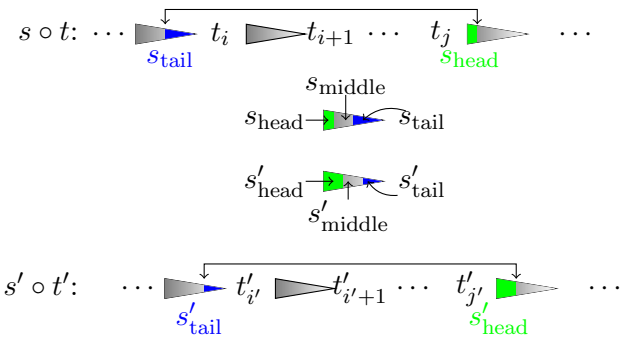


Since $s \sim s'$, there is a bijection $f_{s,s'}$ that maps every substring of s to a substring of s' with the same composition, and a similar bijection $f_{t,t'}$. Let $f_{t,t'}$ map the substring $t_i \cdots t_j$ to a substring $t'_{i'} \cdots t'_{j'}$ of the same composition, and hence length. Since $s \sim s'$, the strings $t_i\, s\, t_{i+1}\, s \cdots s\, t_j$ and $t'_{i'}\, s'\, t'_{i'+1}\, s' \cdots s'\, t'_{j'}$ have the same composition.



Using the tail and head notation above, it remains to show that ℓ and h can be bijectively mapped to ℓ' and h' such that $s_{\text{tail}} s_{\text{head}}$ and $s'_{\text{tail}} s'_{\text{head}}$ have the same composition. There are two cases.

If the length of $s_{\text{tail}} s_{\text{head}}$ is at most n, removing s_{tail} and s_{head} from s yields $s_{\text{middle}} = s_{h+1} \cdots s_{\ell-1}$, which $f_{s,s'}$ maps to a substring $s'_{\text{middle}} = s'_{h'+1} \cdots s'_{\ell'-1}$ of s'. Removing s'_{middle} from s' yields a tail $s'_{\text{tail}} = s'_{\ell'} \cdots s'_n$ and a head $s'_{\text{head}} = s'_1 \cdots s'_{h'}$ of s'. Their combined composition equals that of the original head and tail.



Similarly, if the length of $s_{\text{tail}} s_{\text{head}}$ is greater than n, s_{tail} and s_{head} have a common substring $s_{\text{middle}} = s_\ell \cdots s_h$, which $f_{s,s'}$ maps to a substring $s'_{\text{middle}} = s'_{\ell'} \cdots s_{h'}$, in s'. The head $s'_{\text{head}} = s'_1 \cdots s'_{h'}$ and tail $s'_{\ell'} \cdots s'_n$ of s' have the same combined composition as the original head and tail.

The second generalization of the interleaved construction is to interleaving more than two strings. The following are some properties of interleaving strings.

Associativity: $(s_1 \circ s_2) \circ s_3 = s_1 \circ (s_2 \circ s_3)$.

Unique factorization: Every string has a unique maximal factorization into interleaved strings.

Proof. We first show that if $s \circ s' = t \circ t'$ where s and t are irreducible under the interleaving operation, then both equal $u \circ u'$ for some string u of length $(|s+1|, |t+1|) - 1$. If $|s| = |t|$ then we easily get $s = t$, else from lengths we have $|s+1| \cdot |s'+1| = |t+1| \cdot |t'+1|$ and hence $(|s+1|, |t+1|) = |u+1| > 1$. We can expand $s = u_1 a_1 u_2 a_2 \cdots u_M$, where $M = |s+1|/|u+1|$ and $t = v_1 b_1 v_2 b_2 \cdots v_N$, where $N = |t+1|/|u+1|$ where $|u_i| = |v_i| = |u|$ and each a_i and b_i are bits. Then since $(M, N) = 1$ for any $(i, j) \in ([M], [N])$ there are substrings s_1 and t_1 of $s \circ s'$ and $t \circ t'$ starting and ending at same index such that $s_1 = u_i$ and $t_1 = v_j$. This implies that $u_i = v_j$ for all i and j and thus each of s and t are reducible. □

Theorem 21.

$$E_{s_1 \circ s_2 \circ \cdots \circ s_k} = \{t_1 \circ t_2 \circ \cdots \circ t_k : t_i \sim s_i \text{ for all } i\}.$$

Proof. Proving $\supseteq$ is simple. By associativity and induction, if $s_i \sim t_i$ for $i = 1, \ldots, m$, then

$$s_1 \circ s_2 \circ \cdots \circ s_m \sim t_1 \circ t_2 \circ \cdots \circ t_m.$$

The proof of $\subseteq$ is more complex and postponed to Appendix B. □

The simple ($\supseteq$) part of the theorem suggests a simple construction of large confusable sets. Every string is equicomposable with its reversal, hence

$$E_{s_1 \circ s_2 \circ \cdots \circ s_k} \supseteq \{\tilde{s}_1 \circ \tilde{s}_2 \circ \cdots \circ \tilde{s}_k : \tilde{s}_i \text{ is either } s_i \text{ or } s_i^*\}.$$

If all these strings are non-palindromic, then by unique factorization, each of the resulting interleaved products is different, hence

$$|E_{s_1 \circ s_2 \circ \cdots \circ s_k}| \geq 2^k.$$

Let $|s_i| = m_i$. We saw that $|s_1 \circ s_2| = (m_1 + 1)(m_2 + 1) - 1$, and by induction

$$|s_1 \circ s_2 \circ \cdots \circ s_k| = \prod_{i=1}^{k} (m_i + 1) - 1.$$

For example, taking $s_1 = \ldots = s_k = 01$ and $n = 3^k - 1$,

$$E_n = E_{3^k - 1} \geq |E_{\underbrace{01 \circ 01 \circ \cdots \circ 01}_{k}}| \geq 2^k = (n+1)^{\log_3 2}.$$

For general lengths n, consider the distinct-prime factorization

$$n + 1 = 2^{e_0} p_1^{e_1} p_2^{e_2} \ldots p_k^{e_k}.$$

For every prime $p \geq 3$, take a non-palindromic string of length $p - 1$, and for every pair of 2's, take a non-palindromic string of length $2 \cdot 2 - 1 = 3$. Interleaving all these strings and their reversals lower bounds the largest number of mutually equicomposable strings.

Theorem 22.

$$E_n \geq 2^{\lfloor \frac{e_0}{2} \rfloor + e_1 + e_2 + \cdots + e_k}.$$

As mentioned above, this is the strongest lower bound we have. In the rest of the section we discuss the structure of possible extensions.

The confusable strings we have seen so far were interleaved as in Theorem 21. Exhaustive search shows that so are all confusable strings of length ≤ 22. Yet Theorem 21 shows that equicomposable sets larger than implied by Theorem 22 must be based on non-interleaved strings. We next show examples of non-interleaved confusable strings. However, while showing that non-interleaved confusions may arise, these examples do not necessarily result in larger confusable sets.

We begin with a polynomial interpretation of interleaving.

Lemma 23. *For any s with a 0's and b 1's and any t,*

$$P_{s\circ t}(x,y) = P_s(x,y)P_t(x^{a+1}y^b, x^a y^{b+1}).$$

Proof. Let $t \stackrel{\text{def}}{=} t_1 t_2 \cdots t_m$, so that $s \circ t = s t_1 s t_2 \cdots t_m s$, and define $a_i = \sum_{j=1}^{i} t_j$ and $b_i = \sum_{j=1}^{i}(1-t_j)$ to be the numbers of 0's and 1's in $t_1 t_2 \cdots t_i$. Then,

$$\begin{aligned} P_{s\circ t}(x,y) &= P_s(x,y) + x^{a+t_1}y^{b+(1-t_1)}P_s(x,y) + \cdots + x^{ma+t_1+\cdots+t_m}y^{mb+(1-t_1)+\cdots+(1-t_m)}P_s(x,y) \\ &= P_s(x,y)\sum_{i=1}^{m} x^{ia+a_i}y^{ib+b_i}. \end{aligned}$$

The lemma follows as

$$\sum_{i=1}^{m} x^{ia+a_i}y^{ib+b_i} = \sum_{i=1}^{m} x^{(a_i+b_i)a+a_i}y^{(a_i+b_i)b+b_i} = \sum_{i=1}^{m}(x^{a+1}y^b)^{a_i}(x^a y^{b+1})^{b_i} = P_t(x^{a+1}y^b, x^a y^{b+1}). \quad \square$$

Following are more general classes of equicomposable strings.

The following result follows easily from Theorem 8. This is the most general class of strings which are equicomposable with each other we have come up with.

Theorem 24. *Let $s_1, s_2, \ldots, s_k$ be strings whose generating polynomials have a common factor $D(x,y)$. Derive $t_1, t_2, \ldots, t_k$ from $s_1, s_2, \ldots, s_k$ by replacing $D(x,y)$ by its reciprocal. Then for any $x_1 x_2 \ldots x_{k-1}$,*

$$s_1\, x_1\, s_2\, x_2\, \cdots\, x_{k-1}\, s_k \sim t_1\, x_1\, t_2\, x_2\, \cdots\, x_{k-1}\, t_k. \tag{4}$$

The equicomposable strings mentioned earlier are a special case of the theorem. The following corollaries yield classes of equicomposable strings with some structure.

Corollary 25. *Let $A \sim B$ and for $i = 1, 2, \ldots, k$ let $(s_i, t_i) \in \{(A, B'), (B, A')\}$. Then, for any string $\bar{x} = x_1 x_2 \cdots x_{k-1}$, the two strings in (4) are equicomposable.* $\square$

Corollary 26. *Let s_1, ..., s_k be strings of the same composition. Then for any string s_0 and bits $x_1 x_2 \ldots x_{k-1}$,*

$$(s_1 \circ s_0)x_1(s_2 \circ s_0)x_2 \ldots x_{k-1}(s_k \circ s_0) \sim (s_1 \circ s_0^*)x_1(s_2 \circ s_0^*)x_2 \ldots x_{k-1}(s_k \circ s_0^*). \quad \square$$

For example, taking $s_1 = 010$, $s_2 = 001$, $s_0 = 01$, and $x_1 = 0$, we obtain the 23-bit confusion

$$\begin{aligned} \mathbf{0100101010}\mathit{0}\mathbf{0010011001} &= (\mathbf{010} \circ 01)\mathit{0}(\mathbf{001} \circ 01) \\ &\sim (\mathbf{010} \circ 10)\mathit{0}(\mathbf{001} \circ 10) \\ &= \mathbf{0101010010}\mathit{0}\mathbf{0011001}0\mathbf{001}. \end{aligned}$$

These strings are clearly non-interleaved, hence are the shortest non-interleaved confusable strings.

This example shows in particular that string equicomposability differs from partition coarsening where [21] used Ribbon Schur Functions to show that all confusions arise from interleaving.

11 Upper bounds on E_n

To upper bound the largest number of mutually-equicomposable strings, we use well-known results on the factorization of $x^n - 1$.

Theorem 27. *For all n,*

$$E_n \leq \min\left\{2^{d(n+1)-1}, (n+1)^{1.23}\right\}.$$

Proof. The $(n+1)^{1.23}$ term follows from results in [22] by replacing y in $P(x,y)$ by x^{n+1} to obtain a univariate polynomial, and reproving theorems therein for this polynomial, showing that $E_n \leq (n+1)^{1.23}$.

To prove the first upper bound, let s have length n, then

$$P_s(x,x) = 1 + x + x^2 + x^3 + \ldots + x^n = \frac{x^{n+1}-1}{x-1}.$$

Hence the prime factorization of $P_s(x,x)$ contains $d(n+1)-1$ terms. From Lemma 16 the prime factorization of $P_s(x,y)$ contains at most that many terms, and the bound follows from Corollary 9. □

To also relate the form of this upper bound to that of the lower bound in Theorem 22, again let $n+1 = 2^{e_0}p_1^{e_1}p_2^{e_2}\ldots p_k^{e_k}$ be a factorization into distinct primes. Then $d(n+1) = (e_0+1)(e_1+1)\cdots(e_k+1)$, hence

$$E_n \leq \min\left\{2^{(e_0+1)(e_1+1)\cdots(e_k+1)-1}, (n+1)^{1.23}\right\}.$$

The lower bound in Theorem 22 is tight when $n+1$ is a prime power or twice a prime power.

Theorem 28. *For any k,*

$$E_{2^k-1} = 2^{\lfloor \frac{k}{2} \rfloor},$$

and for prime $p \geq 3$,

$$E_{p^k-1} = 2^k.$$

Proof. Eisenstein's criterion again implies the prime factorization

$$P_s(x,x) = 1 + x + x^2 + \cdots + x^{2^k-1} = \prod_{j=1}^{k-1}(1 + x^{2^j}).$$

Once again, to prove the bound it suffices to show that any factorization of $P_s(x,y)$ into $f(x,y)g(x,y)$ with $g(x,x) = 1 + x^{2^i}$ implies that $g(x,y)$ is palindromic.

Let $s = s_1 s_2 \ldots s_{2^k-1}$ and as before, let a_i be the number of zeros among $s_1, \ldots, s_i$. Then $P(x,y) \stackrel{\text{def}}{=} P_s(x,y) = \sum_{i=0}^{2^k-1} x^{a_i} y^{i-a_i}$. Denote $\dfrac{\partial P}{\partial x}(x,y)|_{y=x}$ by P'_x, then

$$(1-x)P'_x = (1-x)f(x,x)g'_x + (1-x)g(x,x)f'_x$$

$$\implies \sum_{i=0}^{2^k-2} s_{i+1}x^i - a_{2^k-1}x^{2^k-1} = \left(1 - x^{2^i} + x^{2\cdot 2^i} - x^{3\cdot 2^i} + \cdots - x^{(2^{k-i}-1)\cdot 2^i}\right) g'_x + (1+x^{2^i})(1-x)f'_x.$$

The degree of g'_x is $2^i - 1$, and that of $(1-x)f'_x$ is $2^k - 1 - 2^i$, so we can write

$$(1-x)f'_x = \sum_{j=0}^{2^{k-i}-2} x^{j\cdot 2^i} f_j,$$

where each f_j has degree $\leq 2^i - 1$. Defining $f_{-1} = f_{2^k} = 0$ and $s_{2^k} = -a_{2^k-1}$,

$$\sum_{i=0}^{2^k-2} s_{i+1}x^i - a_{2^k-1}x^{2^k-1} = \sum_{j=0}^{2^{k-i}-1} (-1)^j x^{j\cdot 2^i} g'_x + \sum_{j=0}^{2^{k-i}-2} x^{j\cdot 2^i} f_j + \sum_{j=1}^{2^{k-i}-1} x^{j\cdot 2^i} f_j$$

$$\implies \sum_{j=0}^{2^{k-i}-1} x^{j\cdot 2^i} \left(\sum_{l=0}^{2^i-1} s_{j\cdot 2^i+l+1} x^l \right) = \sum_{j=0}^{2^{k-i}-1} x^{j\cdot 2^i} \left((-1)^j g'_x + f_{j-1} + f_j \right)$$

In the above equation, if we sum all the terms with even values of j and subtract from the sum of odd values we get,

$$\sum_{j=0}^{2^{k-i}-1} (-1)^{j\cdot 2^i} \left(\sum_{l=0}^{2^i-1} s_{j\cdot 2^i+l+1} x^l \right) = 2^{k-i} g'_x.$$

Since all s_l's except s_{2^k} are either 0 or 1, we must have

$$g'_x = a \cdot x^{2^i-1},$$

and thus,

$$g(x,y) = 1 + g_1(x,y) + x^a y^{2^i-a}.$$

Claim 1. *If $f(x,x) = K$, a constant and $f'_x = 0$, then $(x-y)^2 | f(x,y) - K$.*

Proof. Let $f(x,y) = (x-y)f_1(x,y) + f_2(x)$. Then, clearly $f_2(x) = K$. Also, $f'_x = 0$ implies $f_1(x,x) = 0$ and thus $f_1(x,y)$ has $x - y$ as a factor. □

Using the claim we conclude that $g(x,y) = 1 + (x-y)^2 g_1(x,y) + x^a y^{2^i-a}$. Consider,

$$f(x,y)\,(g - g^*)\,.$$

It is the difference of two generating polynomials, but substituting $y = -x$, we see that each coefficient is divisible by 4 which cannot be true unless $g - g^* = 0$.

It is easy to see that the bound is achievable by taking strings of the form $t_0 \circ t_1 \circ t_2 \circ \cdots \circ t_{\lfloor \frac{k}{2} \rfloor}$ where $t_i = 001$ for $i \geq 1$, and t_0 is the empty string when k is even, else it is the single bit 0.

When $n = p^k - 1$ for a prime $p \geq 3$, Theorems 22 and 27 coincide. □

12 Reconstruction algorithm

We present an algorithm for reconstructing strings from their composition multiset. The algorithm takes as input the composition multiset $\mathcal{S}_s$ of a string s over an alphabet Σ, and elements of E_s, the set of all strings confusable with s. We show that for alphabet size ≥ 4, the string s that generated $\mathcal{S}_s$ is added in quadratic time. The algorithm successively reconstructs s from both ends and backtracks when it errs. It can be viewed as a modification of a similar algorithm for the turnpike problem.

We first establish two properties of $\mathcal{S}_s$ that help reduce the algorithm's search space. The next lemma shows that the composition multiset determines the set $\{s_i, s_{n+1-i}\}$ of symbols at the symmetric positions i and $(n+1-i)$ for $i = 1, 2, \ldots, \lceil \frac{n}{2} \rceil$.

Lemma 29. *$\mathcal{S}_s$ determines $\{s_i, s_{n+1-i}\}$ for $i = 1, 2, \ldots, \lceil \frac{n}{2} \rceil$.*

Proof. Let the union of compositions be their union as multisets. For example, $A^2B \cup ABC^2 \cup AC = A^4B^2C^3$. For a string s, let M_i denote the union of the compositions of all substrings of length i. For example, for ABAC, $M_1 = A^2BC$, $M_2 = A^3B^2C$. Note that all M_i's can be easily determined from the string, and that for $1 \le i \le \lfloor n/2 \rfloor$, $M_{n+1-i} = M_i$. For a multiset S, let $j \cdot S$ be the j-fold union $S \cup \ldots \cup S$. It is easy to see that

$$M_2 \cup \{s_1, s_n\} = 2 \cdot M_1,$$

hence $\{s_1, s_n\}$, can be deduced from $\mathcal{S}_s$. More generally, for $i = 2, \ldots, \lceil \frac{n}{2} \rceil$,

$$M_i \cup \{s_{i-1}, s_{n-i+2}\} \cup 2 \cdot \{s_{i-2}, s_{n-i+3}\} \cup \ldots \cup (i-1) \cdot \{s_1, s_n\} = i \cdot M_1.$$

Using this equation inductively over $i = 2, \ldots, \lceil \frac{n}{2} \rceil$, yields all multisets $\{s_i, s_{n-i+1}\}$. □

For $1 \le i < \lceil n/2 \rceil$, let $\mathcal{T}_i$ be the collection of compositions of strings s_j^k where $j, k \le i$, or $j, k \ge n+1-i$, or $j \le i+1 \le n-i \le k$, namely the collection of compositions of all strings that are either on "one side" of s_{i+1}^{n-i} or "straddle" s_{i+1}^{n-i}. The next lemma shows that the composition of the whole string, along with the strings s^i and s_{n+1-i}^n determine $\mathcal{T}_i$.

Lemma 30. *$\mathcal{S}_s$, s^i, and s_{n+1-i}^n determine $\mathcal{T}_i$.*

Proof. $\mathcal{T}_i$ consists of compositions of three types of strings, those that are substrings of s_1^i, that are substrings of s_{n+1-i}^n, and substrings that cover all symbols in between, *i.e.,* s_{i+1}^{n-i}. The first and last i symbols determine the compositions of the first two type of strings. The third type of strings are those that contain the string s_{i+1}^{n-i}. This is a *symmetric substring*, namely it has the same number of symbols on its left and right, and by Lemma 29 we can determine its composition. Knowing this composition and the first and last i symbols yields the multiset of strings of the third type. □

We use the two lemmas to reconstruct the string. Let $w(s)$ denote the composition of a string s.

Lemma 31. *If $w(s_1^i) \ne w(s_{n-i+1}^n)$, then $\mathcal{S}_s$, s_1^i, and s_{n+1-i}^n determine the pair (s_{i+1}, s_{n-i}).*

Proof. By Lemma 30, we can determine $\mathcal{T}_i$. Consider the two longest compositions in $\mathcal{S}_s \setminus \mathcal{T}_i$. They correspond to the length-$(n-i-1)$ strings s_1^{n-i-1} and s_{i+2}^n. The complements of these two compositions are therefore the compositions of s^{i+1} and s_{n-i}^n.

By Lemma 29, we can also derive the multiset $\{s_{i+1}, s_{n-i}\}$. If $s_{i+1} = s_{n-i}$, then we can determine s^{i+1} and s_{n-i}^n. Otherwise, $s_{i+1} \ne s_{n-i}$, and since $w(s_1^i) \ne w(s_{n+1-i}^n)$, it is easy to see that $\{w(s_1^i) \cup w(s_{i+1}), w(s_{n-i}) \cup w(s_{n-i+1}^n)\} \ne \{w(s_1^i) \cup w(s_{n-i}), w(s_{i+1}) \cup w(s_{n-i+1}^n)\}$, hence we can determine the pair (s_{i+1}, s_{n-i}). □

The algorithm reconstructs the string by sequentially deciding on the values of the pair of symbols s_i and s_{n-i+1}, that are in symmetric positions. It starts with the symbols s_1 and s_n at the ends of the string, and progressively moves closer to the center. The algorithm first determines $s_1 s_2$ and $s_{n-1} s_n$, uniquely upto reversal. By Lemma 29, we know the multiset $\{s_1, s_n\}$ and can arbitrarily decide which is s_1 and which is s_n. It next determines s_2 and s_{n-1}. Again by the lemma we know $\{s_2, s_{n-1}\}$, and if $s_1 = s_n$ we can decide on s_2 and s_{n-1} arbitrarily, while if $s_1 \ne s_n$, by Lemma 31, we can determine s_2 and s_{n-1}. For $\{s_3, s_{n-2}\}$, the ends s_1^2 and s_{n-1}^n can differ, while their weights could be the same. In such cases Lemma 31 cannot be applied. However if $s_3 = s_{n-2}$, which can be determined from the Lemma 29, we easily determine (s_3, s_{n-2}). In other words, if $s_i = s_{n+1-i}$, the algorithm easily determines (s_i, s_{n+1-i}). Therefore, from this point on, when the algorithm has reconstructed the first and last i symbols, and $w(s_1^i) = w(s_{n+1-i}^n)$ but $s_{i+1} \ne s_{n+1-i}$, it guesses one of the two possibilities for (s_{i+1}, s_{n+1-i}) and proceeds, while keeping track of the number t and locations

$i_1 < i_2 < \ldots < i_t$ of locations where guesses were made. After determining/guessing s_{i+1} and s_{n-i}, the algorithm updates $\mathcal{T}_i$ to T_{i+1}, which can be accomplished in near-linear time. It then checks whether $\mathcal{T}_{i+1} \subseteq \mathcal{S}_s$ as multisets. If at some point $\mathcal{T}_{i+1} \subsetneq \mathcal{S}_s$, the algorithm *backtracks*. It changes its guess at location i_t by swapping s_{i_t} and s_{n+1-i_t}, changes t to $t-1$, and restarts reconstruction from location i_t+1. The process continues until the whole string is reconstructed, namely $i = \lceil \frac{n}{2} \rceil$ and $\mathcal{T}_{\lceil n/2 \rceil} = \mathcal{S}_s$. If at that point there are $t \geq 1$ guesses, then as before the algorithm backtracks to guess t and tries to find additional strings in E_s.

Since our algorithm relies on the compositions of strings from both ends, it is helpful to define a few terms. Let

$$\ell_s \stackrel{\text{def}}{=} \left|\{i < n/2 : w(s^i) = w(s_{n+1-i}^n) \text{ and } s_{i+1} \neq s_{n-i}\}\right|$$

be the number of substrings from the ends having the same composition, and the next two symbols are distinct. Let

$$L_s \stackrel{\text{def}}{=} \max_{t \in E_s} \ell_t$$

be the largest value of ℓ over all strings in E_s.

The backtracking algorithm induces a binary tree where the nodes represent locations at which there are two possible reconstructions. The procedure described above does a depth-first search. Instead we could also do a breadth first search, where we consider all branches at once, namely at any time, all potential reconstructions correspond to level t or $t+1$. This implies that given $\mathcal{S}_s$, the algorithm is able to find s before depth $\ell_s + 1$.

We bound the time required to reconstruct s from $\mathcal{S}_s$ in the following theorem.

Theorem 32. *The backtracking algorithm, run using proper data structures, given $\mathcal{S}_s$ and ℓ_s, outputs a set of strings that contains s in time $O(2^{\ell_s} n^2 \log n)$. Futhermore, E_s can be recovered in time $O(2^{L_s} n^2 \log n)$.*

Proof. We assume an arbitrary order over the elements of Σ. This introduces a lexicographical ordering over compositions of strings on Σ. We use a Red-Black tree [23] to store multisets of compositions. The advantage of this data structure is that insertion, deletion and search all require time $O(\log n)$, where n is the size of the data.

Notice that $\mathcal{T}_{i+1}$ is obtained from $\mathcal{T}_i$ by adding the compositions of substrings that have an endpoint at s_{i+1} or s_{n-i}. In particular, at most $2n$ compositions are added, requiring $O(n \log n)$ time. For each branch, we keep a copy of $\mathcal{S}_s$ and we prune it to populate $\mathcal{T}$, *i.e.,* while constructing $\mathcal{T}_{i+1}$ we simultaneously remove the new entries/compositions from the copy of $\mathcal{S}$ corresponding to this branch, requiring another $O(n \log n)$. When there are two possibilities for reconstruction at any stage, we make copies of $\mathcal{T}_i$ and $\mathcal{S}$ corresponding to that node and proceed along each. This step takes time $O(n^2 \log n)$.

The algorithm while reconstructing s does not "fork out" more then $\ell_s + 1$ times, therefore the number of branches before reconstructing s is at most 2^{ℓ_s}. Along each path we make $n/2$ iterations requiring a total time of $O(n^2 \log n)$. Combining these, the total complexity of reconstructing s is at most $O(2^{\ell_s} n^2 \log n)$.

To reconstruct E_s, we note that the number of "forks" is $\leq L_s$. Therefore, a similar computation shows that the algorithm runs in time $O(2^{L_s} n^2 \log n)$. □

We bounded the run time of reconstructing a string as a function of ℓ_s. For random strings, we bound the run time by studying the distribution of ℓ_s.

We do this by using some well known results in random walk theory. These are applications of the Striling approximation of factorials.

Lemma 33 (Stirling's approximation)**.** *For any $n \geq 1$, there is a $\theta_n \in (\frac{1}{12n+1}, \frac{1}{12n})$ such that*

$$n! = \sqrt{2\pi n} \left(\frac{n}{e}\right)^n e^{\theta_n}.$$

We now apply this to bound the probability that two random strings have the same composition. A stronger version of the result that finds asymptotic equality is proved in [24], but is not required for our purposes.

Lemma 34. *For $|\Sigma| = k$, let s and t be two random length-n strings over Σ, then for $k \geq n$*

$$P\left(w(s) = w(t)\right) < \frac{n!}{k^n},$$

and for $k < n$

$$P\left(w(s) = w(t)\right) < \frac{k^{k/2} e^{1/12n}}{(2\pi n)^{(k-1)/2}}.$$

Proof. The probability that the symbols in Σ appear $i_1, \ldots, i_k$ times respectively in a random length-n string is

$$\frac{1}{k^n} \binom{n}{i_1, \ldots, i_k}.$$

Therefore, the probability that two random strings have the same composition is

$$\begin{aligned}\sum_{i_1+\ldots+i_k=n} \frac{1}{k^{2n}} \binom{n}{i_1, \ldots, i_k}^2 \leq & \frac{1}{k^n} \max_{i_1+\ldots+i_k=n} \binom{n}{i_1 \ldots i_k} \cdot \sum_{i_1+\ldots+i_k=n} \frac{1}{k^n} \binom{n}{i_1, \ldots, i_k} \\ = & \frac{1}{k^n} \max_{i_1,\ldots,i_k} \binom{n}{i_1 \ldots i_k},\end{aligned}$$

where the last step follows from

$$\sum_{i_1+\ldots+i_k=n} \frac{1}{k^n} \binom{n}{i_1, \ldots, i_k} = 1.$$

For $k \geq n$,

$$\max_{i_1,\ldots,i_k} \binom{n}{i_1 \ldots i_k} = n!$$

implying the first part.

For $n \geq k$, note that

$$f(x) \stackrel{\text{def}}{=} \sqrt{2\pi x} \left(\frac{x}{e}\right)^x$$

is convex in $(1, \infty)$. Therefore,

$$\prod_{j=1}^{k} i_j! \stackrel{(a)}{>} \prod_{j=1}^{k} f(i_j) \stackrel{(b)}{>} \left(f\left(\frac{n}{k}\right)\right)^k = \left(\sqrt{2\pi \frac{n}{k}}\right)^k \left(\frac{n}{ke}\right)^n.$$

where (a) uses Stirling's approximation, and (b) follows from convexity of f. Hence

$$\max_{i_1,\ldots,i_k} \binom{n}{i_1 \ldots i_k} < \frac{n!}{\left(f\left(\frac{n}{k}\right)\right)^k} < \frac{k^{k/2} e^{1/12n}}{(2\pi n)^{(k-1)/2}},$$

where in the last step we approximate $n!$. □

We now study the distribution of ℓ_s. We first consider alphabet size ≥ 4 in detail, and then provide performance guarantees for alphabet size 3 and 2.

13 Alphabet size ≥ 4

Consider two uniformly random independent infinite strings s^∞ and t^∞ over Σ. A set of integers $i_1 < i_2 < \ldots < i_m$ is *non-consecutive* if $i_{j+1} > i_j + 1$. Let F_m be the event that there are at least $m+1$ *non-consecutive* integers $i_0 \stackrel{\text{def}}{=} 0 < i_1 < i_2 < \ldots < i_m$ such that for each $j \leq m$, $w(s_1^{i_j}) = w(t_1^{i_j})$. After a location i_j at which $w(s_1^{i_j}) = w(t_1^{i_j})$ by independence the process is equivalent to starting at time 0. It follows that $P(F_{j+1}|F_j) = P(F_1)$. Therefore,

$$P(F_m) = P(F_1)^m. \tag{5}$$

Let $M(s^\infty, t^\infty)$ denote the total number of non-consecutive integers for which $w(s_1^i) = w(t_1^i)$. Then by Equation (5),

$$\mathbb{E}[M] = \sum_{m \geq 1} P(M \geq m) = \sum_{m \geq 1} P(F_m) = \frac{P(F_1)}{1 - P(F_1)}.$$

However, if instead of non-consecutiveness, we only restrict $i_1 \geq 2$, Lemma 34 shows that

$$\mathbb{E}[M] \leq \sum_{n=2}^{k} \frac{n!}{k^n} + \sum_{n=k+1}^{\infty} \frac{k^{k/2} e^{1/12n}}{(2\pi n)^{(k-1)/2}}.$$

The right hand side of this equation is finite for $k \geq 4$. In fact, it decays as $O(1/k^2)$ with k. This implies that $p_k \stackrel{\text{def}}{=} P(F_1) < 1$ and therefore for a random string s, Equation (5) gives

$$P(\ell_s > m) \leq P(F_m) = p_k^m,$$

proving the following lemma.

Lemma 35. *A random string over alphabet size $k \geq 4$, with probability $\geq 1 - p_k^m$, satisfies $\ell_s \leq m$.*

This implies that with probability $\geq 1 - \delta$, $\ell_s < \log_{1/p_k} \frac{1}{\delta}$. Therefore applying Theorem 32,

Theorem 36. *For a random string s over an alphabet of size ≥ 4 the backtracking algorithm with probability $> 1 - \delta$, outputs a subset of E_s containing s in time $O_{\delta,k}(n^2 \log n)$.*

Recall from Theorem 27 that $|E_s| < n^{1.23}$ and therefore by the union bound,

Lemma 37. *With probability $\geq 1 - n^{1.23} p_k^m$, a random string over alphabet size $k \geq 4$ satisfies $L_s < m$.*

Applying these two lemmas to Theorem 32 we get

Theorem 38. *For a random string s over an alphabet of size ≥ 4 the backtracking algorithm with probability $> 1 - \delta$ outputs E_s in time $O_\delta\left(n^{1.23 \log_{1/p_k} 2} n^2 \log n\right)$.*

Recall that p_k falls as $1/k^2$, and with growing k, the algorithm reconstructs entire E_s in near quadratic time. We now consider strings over alphabet sizes 3 and 2.

14 Alphabet size 3 and 2

Recall that the harmonic sum $H_n = 1 + 1/2 + \ldots + 1/n$ converges to $\ln n + \gamma$, where γ is the Euler-Mascheroni constant. Therefore, by Lemma 34 we can bound

$$\mathbb{E}[\ell_s] < 0.84 \ln n + 0.84\gamma.$$

and Markov's inequality with probability $\geq 1 - \delta$, a random string over alphabet of size 3 satisfies $\ell_s < \frac{0.84 \ln n + 0.84\gamma}{\delta}$. Again applying Theorem 32,

Theorem 39. *The backtracking algorithm for random strings over alphabet size 3 outputs a subset of E_s containing s in time $O\left(n^{\frac{0.6}{\delta}}\right)$.*

For alphabet size 2, Lemma 34 shows that

$$\mathbb{E}[\ell_s] < 1.9\sqrt{n} + 0.84\gamma.$$

Using Markov's inequality and Theorem 32 as done for alphabet size 3 yields

Theorem 40. *The backtracking algorithm for random strings over alphabet size 2 outputs a subset of E_s containing s in time $O\left(2^{\frac{2\sqrt{n}}{\delta}}\right)$.*

15 Conclusions and extensions

Starting with protein mass-spectrometry reconstruction, we made two simplifying assumptions: that all bond cuts are equally likely, and that substring weights imply their compositions. These two assumptions reduced protein-reconstruction to the simple problem of reconstructing a string from its substring compositions. We noted that this is the only unstudied variation of four related reconstruction problems, that solving it for binary strings suffices, and that it is also a combinatorial simplification of the long-open turnpike problem.

We called strings with the same composition multiset *equicomposable*, strings equicomposable only with themselves and their reversal *reconstructable*, and those with more than these two trivial equicomposable strings, *confusable*. We noted that all strings of length at most seven are reconstructable, interleaved strings and their reversals to obtain confusable 8-bit strings, and obtained similar confusable strings of all lengths one short of a product of two integers ≥ 3.

Extending polynomial techniques used for the turnpike problem, we represented strings as bivariate *generating* polynomials. We used this formulation to characterize equicomposability in terms of both polynomial multiplication and polynomial factorization, showing in particular that equicomposable strings are determined exactly by the prime factorization of their generating polynomials. We then showed that all strings of lengths not included in the earlier construction, namely, seven, and one less than either a prime or twice a prime, are reconstructable.

Interleaving multiple strings we constructed sets of $(n+1)^{\log_3 2}$ mutually-equicomposable n-bit strings, and showed a pair of non-interleaved confusable strings. Using cyclotomic polynomials we upper bounded the largest number of confusable strings, showing in particular that the lower bound is tight when the sequence length is one short of a prime power and twice a prime power.

Many questions remain. All confusable strings we are aware of are described by Theorem 24. Finding other confusable strings, or proving that this describes all confusable strings would be of interest. We made two assumptions. The first implied that we are given the composition of all substrings. What happens when we are given a fraction of all substrings, or of all substrings of a given length? The proofs provided here use algebraic arguments. Direct combinatorial proofs would be of interest. The second assumption was that all compositions are given correctly. It would be interesting to know whether some errors can be tolerated.

While prime-related reconstructability may be interesting, reconstructability for arbitrary lengths may be of more practical relevance. It would be interesting to determine whether the lower bound of the size of equicomposable sets in Theorem 22 is always tight. It would mean that every string is confusable with at most a sublinear number of strings. A related question is whether most strings of a given length are reconstructable. This question is related to the open problem of whether most 0-1 polynomials are irreducible over the integers [25].A related question addresses the number of composition multisets. If this number is close to 2^n then most strings can be reconstructed. Another variation is when instead of a string, the bits are

arranged on a ring. The constructions presented here extend to ring. Proving the upper bounds is still open.

Other problems relate to algorithms for reconstructing a string from its substring composition. As noted earlier, n-bit reconstruction can be reduced to solving a turnpike problem with $n+1$ exits, and total length $\leq n^2$. This implies a polynomial-time algorithm for the reconstruction. However such a generic algorithm may have high complexity and an algorithm that uses the structure of the reconstruction problem is of interest.

Acknowledgements

We thank Sampath Kannan, Ananda Theertha Suresh and Alex Vardy for helpful discussions and suggestions.

References

[1] T. E. Creighton, *Proteins: Structures and Molecular Properties*, 2nd ed. W. H. Freeman, 1992.

[2] D. W. Mount, *Bioinformatics: Sequence and Genome Analysis*, 2nd ed. Cold Spring Harbor Laboratory Press, 2001.

[3] T. Dakic, "On the turnpike problem," Ph.D. dissertation, Simon Fraser University, 2000.

[4] S. S. Skiena, W. D. Smith, and P. Lemke, "Reconstructing sets from interpoint distances (extended abstract)," in *Symposium on Computational Geometry*, 1990, pp. 332–339.

[5] D. Margaritis and S. S. Skiena, "Reconstructing strings from substrings in rounds," in *Proceedings of the 36th Annual Symposium on Foundations of Computer Science*, 1995, pp. 613–620.

[6] V. I. Levenshtein, "Efficient reconstruction of sequences from their subsequences or supersequences," *Journal of Combinatorial Theory, Series A*, vol. 93, no. 2, pp. 310–332, 2001.

[7] M. Dudik and L. J. Schulman, "Reconstruction from subsequences," *Journal of Combinatorial Theory, Series A*, vol. 103, no. 2, pp. 337–348, 2003.

[8] T. Batu, S. Kannan, S. Khanna, and A. McGregor, "Reconstructing strings from random traces," in *Proceedings of the fifteenth annual ACM-SIAM symposium on Discrete algorithms*, 2004, pp. 910–918.

[9] K. Viswanathan and R. Swaminathan, "Improved string reconstruction over insertion-deletion channels," in *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, 2008, pp. 399–408.

[10] J. Acharya, H. Das, O. Milenkovic, A. Orlitsky, and S. Pan, "On reconstructing a string from its substring compositions," in *Proceedings of IEEE Symposium on Information Theory*, 2010.

[11] A. L. Patterson, "A direct method for the determination of the components of interatomie distanees in crystals," *Zeitsehr. Krist.*, vol. 90, pp. 517–554, 1935.

[12] ——, "Ambiguities in the X-ray analysis of crystal strycture," *Phys. Review.*, vol. 65, pp. 195–201, 1944.

[13] T. I. Dix and D. H. Kieronska, "Errors between sites in restriction site mapping," *Comput. Appl. Biosci.*, vol. 4, pp. 117–123, 1988.

[14] A. Daurat, Y. Gérard, and M. Nivat, "Some necessary clarifications about the chords' problem and the partial digest problem," *Theor. Comput. Sci.*, vol. 347, no. 1-2, pp. 432–436, 2005.

[15] S. Chen, Z. Huang, and S. Kannan, "Reconstructing numbers from pairwise function values," in *International Symposium on Algorithms and Computation*, 2009, pp. 142–152.

[16] G. Fici and Z. Liptk, "On prefix normal words," in *Developments in Language Theory*, ser. Lecture Notes in Computer Science, G. Mauri and A. Leporati, Eds. Springer Berlin Heidelberg, 2011, vol. 6795, pp. 228–238.

[17] J. Acharya, H. Das, O. Milenkovic, A. Orlitsky, and S. Pan, Private communication with S. Kannan, March 2009.

[18] A. K. Lenstra, H. W. Lenstra, and L. Lovász, "Factoring polynomials with rational coefficients," *Math. Ann.*, vol. 261, pp. 515–534, 1982.

[19] J. Rosenblatt and P. D. Seymour, "The structure of homometric sets," *SIAM Journal on Algebraic and Discrete Methods*, vol. 3, no. 3, pp. 343–350, 1982.

[20] S. Lang, *Algebra*, 3rd ed. Graduate Texts in Mathematics, Springer-Verlag, 2002.

[21] L. J. Billera, H. Thomas, and S. van Willigenburg, "Decomposable compositions, symmetric quasisymmetric functions and equality of ribbon Schur functions," *Advanced Mathematics*, vol. 204, no. 1, pp. 204–240, 2006.

[22] P. Lemke and M. Werman, "Inverting the autocorrelation and the problem of locating points on a line, given unlabelled distances between them," IMA Preprint, Institute for Mathematics and Its Applications, Tech. Rep. 453, 1988.

[23] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*, 2nd ed. The MIT Press, 2001.

[24] L. B. Richmond and J. O. Shallit, "Counting abelian squares," *Electronic Journal of Combinatorics*, vol. 16, no. 1, pp. 317–348, 2009.

[25] A. M. Odlyzko and B. Poonen, "Zeros of polynomials with 0, 1 coefficients," *Enseignement Mathématique*, vol. 39, no. 2, pp. 317–348, 1993.

A Proof of Theorem 17

The proof follows the proof of Theorem 20. We show that in any factorization

$$P_s(x,y) = f(x,y)g(x,y),$$

at least one of $f(x,y)$ and $g(x,y)$ is palindromic. By Lemma 19, s is reconstructable.

$$\begin{aligned} P_s(x,x) &= 1 + x + x^2 + \cdots + x^7 \\ &= (1+x)(1+x^2)(1+x^4). \end{aligned}$$

Hence there are only three factorizations of $P_s(x,y)$ into two factors.

Case 1: $f(x,x) = 1+x$ and $g(x,x) = (1+x^2)(1+x^4)$,
This is identical to *Case 1* of Theorem 20.

Case 2: $f(x,x) = 1+x^2$ and $g(x,x) = (1+x)(1+x^4)$,
For case 2 we note that $f(x,y)$ has no linear terms, otherwise it must have both the terms x and y which would imply the existence of terms of the form x^a and y^b for positive integers a and b which violates **G3**.

Case 3: $f(x,x) = 1+x^4$ and $g(x,x) = (1+x)(1+x^2)$.

In Case 3 concentrate on $g(x,x) = (1+x)(1+x^2) = 1+x+x^2+x^3$. Just as in the case 2 of the proof of the Theorem 20 we can show here that the string is of the form tut where t is a length 3 string. But this implies that $P_s = P_t(1+x^a y^b)$ where a abd b are the number of 0's and 1's in tu, so $a+b=4$. Since $1+x^4$ is irreducible, this factor is palindromic which proves the third case as well.

B Proof of $\subseteq$ in Theorem 21

We first show the following lemma.

Lemma 41. *Let $P(x,y)$ be a generating polynomial. Any $Q(x,y) \in \mathbb{Z}[x,y]$ with constant term 1 satisfying*

$$P(x^{a+1}y^b, x^a y^{b+1}) P^*(x^{a+1}y^b, x^a y^{b+1}) = Q(x,y)Q^*(x,y)$$

has the form $R(x^{a+1}y^b, x^a y^{b+1})$, where $R(x,y)$ is a generating polynomial.

Proof: Similar to the proof of Lemma 5, we can show that

$$P(x^{a+1}y^b, x^a y^{b+1}) P\left(\frac{1}{x^{a+1}y^b}, \frac{1}{x^a y^{b+1}}\right) = Q(x,y)Q\left(\frac{1}{x}, \frac{1}{y}\right), \tag{6}$$

and $Q(x,y)$ is a 0-1 polynomial.

Note that, all terms in the expansion of the left-hand side of Equation (6) has the form

$$\left(x^{a+1}y^b\right)^s \left(x^a y^{b+1}\right)^t \cdot \frac{1}{\left(x^{a+1}y^b\right)^{s'} \left(x^a y^{b+1}\right)^{t'}} = x^{(a+1)(s-s')+a(t-t')}\, y^{b(s-s')+(b+1)(t-t')}.$$

For any term $x^i y^j$ present in $Q(x,y)$, consider terms of the form $x^i y^j \cdot 1$ in the expansion of the right-hand side of Equation (6). Then, for some s,t and s',t',

$$\begin{aligned} i &= (a+1)(s-s') + a(t-t'), \\ j &= b(s-s') + (b+1)(t-t'). \end{aligned}$$

For simplicity, let $u_{ij} = s-s' \in \mathbb{Z}$ and $v_{i,j} = t-t' \in \mathbb{Z}$. Then

$$Q(x,y) = \sum_{i,j} x^i y^j = \sum_{i,j} \left(x^{a+1}y^b\right)^{u_{ij}} \left(x^a y^{b+1}\right)^{v_{ij}}.$$

Let $R(x,y) = \sum_{i,j} x^{u_{ij}} y^{v_{ij}}$. Then $Q(x,y) = R(x^{a+1}y^b, x^a y^{b+1})$.

It remains to show that $R(x,y)$ is a generating polynomial. Let $u \stackrel{\text{def}}{=} \min u_{ij}$, and $v \stackrel{\text{def}}{=} \min v_{ij}$. Then $T(x,y) \stackrel{\text{def}}{=} x^{-u}y^{-v}R(x,y) \in \mathbb{Z}[x,y]$. It's straightforward to see that

$$P(x,y)P\left(\frac{1}{x}, \frac{1}{y}\right) = T(x,y)T\left(\frac{1}{x}, \frac{1}{y}\right).$$

It follows from Lemma 5 that $T(x,y)$ is a generating polynomial. Note that

$$Q(x,y) = (x^{a+1}y^b)^u (x^a y^{b+1})^v \, T\left(x^{a+1}y^b, x^a y^{b+1}\right).$$

Since both $Q(x,y)$ and $T\left(x^{a+1}y^b, x^a y^{b+1}\right)$ are polynomials with constant term 1, we must have $(a+1)u + av = 0$, and $bu + (b+1)v = 0$, which imply that $u = v = 0$. Hence $R(x,y) = T(x,y)$ is a generating function, and

$$Q(x,y) = R(x^{a+1}y^b, x^a y^{b+1}). \qquad \square$$

Proof of Theorem 21: It suffices to consider $k = 2$, i.e., $s = s_1 \circ s_2$. For $i = 1, 2$, let $P_i(x,y)$ be the generating polynomial of S_i. By Lemma 23, the generating polynomial of s is $P_s(x,y) = P_{s_1}(x,y) P_{s_2}\left(x^{a+1}y^b, x^a y^{b+1}\right)$, where a and b are the numbers of ones and zeroes in s_1. Let

$$P_{s_1}(x,y) = \prod_{i=1}^{k_1} A_i(x,y), \text{ and } P_{s_2}\left(x^{a+1}y^b, x^a y^{b+1}\right) = \prod_{i=1}^{k_2} B_i(x,y),$$

where $A_1, A_2, \ldots A_{k_1}$ and $B_1, B_2, \ldots, B_{k_2}$ are irreducible factors. Then

$$P_s(x,y) = \prod_{i=1}^{k_1} A_i(x,y) \prod_{i=1}^{k_2} B_i(x,y).$$

Since $s \sim t$, by Theorem 8, there exist $K_1 \subseteq [k_1]$ and $K_2 \subseteq [k_2]$ such that

$$P_t(x,y) = \prod_{i \in K_1} A_i(x,y) \prod_{i \in [k_1] \setminus K_1} A_i^*(x,y) \prod_{i \in K_2} B_i(x,y) \prod_{i \in [k_2] \setminus K_2} B_i^*(x,y).$$

Let

$$Q_1(x,y) \stackrel{\text{def}}{=} \prod_{i \in K_1} A_i(x,y) \prod_{i \in [k_1] \setminus K_1} A_i^*(x,y),$$
$$Q_2(x,y) \stackrel{\text{def}}{=} \prod_{i \in K_2} B_i(x,y) \prod_{i \in [k_2] \setminus K_2} B_i^*(x,y).$$

Note that $Q_1(x,y)$ and $Q_2(x,y)$ satisfies

$$Q_1(x,y) Q_1^*(x,y) = P_{s_1}(x,y) P_{s_1}^*(x,y),$$
$$Q_2(x,y) Q_2^*(x,y) = P_{s_2}(x^{a+1}y^b, x^a y^{b+1}) P_{s_2}^*(x^{a+1}y^b, x^a y^{b+1}).$$

The first equation and Lemma 5 imply that $Q_1(x,y)$ is a generating polynomial. Since both P_t and Q_1 have constant term 1, $Q_2(x,y)$ also has constant term 1. Then the second equation and Lemma 41 imply that

$$Q_2(x,y) = R(x^{a+1}y^b, x^a y^{b+1}),$$

where $R(x,y)$ is a generating polynomial. By Lemma 23, $t = t_1 \circ t_2$, where t_1 has generating polynomial $Q_1(x,y)$, and t_2 has generating polynomial $R(x,y)$. $\square$